

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:  
**09.03.2005 Bulletin 2005/10**

(51) Int Cl.7: **H04L 29/08, H04L 12/56**

(21) Application number: **04017686.9**

(22) Date of filing: **26.07.2004**

(84) Designated Contracting States:  
**AT BE BG CH CY CZ DE DK EE ES FI FR GB GR**  
**HU IE IT LI LU MC NL PL PT RO SE SI SK TR**  
Designated Extension States:  
**AL HR LT LV MK**

• **Fan, Kan F.**  
**Diamond Bar, CA 91765 (US)**  
• **Lindsay, Steven**  
**Mission Viejo, CA 92692 (US)**  
• **McDaniel, Scott S.**  
**Villa Park, CA 92861 (US)**

(30) Priority: **29.08.2003 US 652270**


(71) Applicant: **Broadcom Corporation**  
**Irvine, California 92618-7013 (US)**

(74) Representative: **Jehle, Volker Armin, Dipl.-Ing.**  
**Bosch, Graf von Stosch, Jehle,**  
**Flüggensstrasse 13**  
**80639 München (DE)**

(72) Inventors:  
• **Elzur, Uri**  
**Irvine, CA 92606 (US)**

(54) **System and method for TCP/IP offload independent of bandwidth delay product**

(57) Aspects of the invention may provide TCP offload, which may include acquiring TCP connection variables from a host and managing at least one TCP connection using the acquired TCP connection variables. At least a portion of the acquired TCP connection variables may be updated and at least some of the updated TCP connection variables may be transferred back to the host. In an aspect of the invention, the TCP connection variables may be variables that are independent of bandwidth delay product. At least a portion of the updated TCP connection variables may be utilized by the host to process the TCP connection or another TCP connection. The host may push the variables onto the stack and the TOE may pull the variables from the stack. Also, updated TCP connection variables may be pushed on the stack by the TOE and pulled from the stack by the host.



**FIG. 1**

## Description

### CROSS-REFERENCE TO RELATED APPLICATIONS/INCORPORATION BY REFERENCE

[0001] This application makes reference to, claims priority to and claims benefit from: United States Provisional Patent Application Serial No. 60/408,617, entitled "System and Method for TCP/IP Offload" filed on September 6, 2002; and United States Provisional Patent Application Serial No. 60/407,165, filed on August 30, 2002.

[0002] The above stated application is incorporated herein by reference in its entirety.

### FIELD OF THE INVENTION

[0003] Certain embodiments of the present invention relate to processing of TCP data and related TCP information. More specifically, certain embodiments relate to a method and system for TCP/IP offload independent of bandwidth delay product.

### BACKGROUND OF THE INVENTION

[0004] The initial development of transmission control protocol (TCP) was based on networking and processing capabilities that were then currently available. As a result, various fundamental assumptions regarding its operation were prefaced on networking and processor technologies that existed at that time. Among the assumptions on which TCP was prefaced includes the scarcity and high cost of bandwidth and the partially limitless processing resources available by a host processor. With the advent of technologies such as Gigabit Ethernet (GbE), these fundamental assumptions have radically changed to the point where bandwidth is no longer as scarce and expensive and the host processing resources are now regarded as being limited rather than virtually infinite. In this regard, the bottleneck has shifted from the network bandwidth to the host processing bandwidth. Since host processing systems do more than merely providing faster network connections, shifting network resources to provide much faster network connections will do little to address the fundamental change in assumptions. Notably, shifting network resources to provide much faster network connections would occur at the expense of executing system applications, thereby resulting in degradation of system performance.

[0005] Although new networking architectures and protocols could be created to address the fundamental shift in assumptions, the new architectures and protocols would still have to provide support for current and legacy systems. Accordingly, solutions are required to address the shift in assumptions and to alleviate any bottlenecks that may result with host processing systems. A transmission control protocol offload engine

(TOE) may be utilized to redistribute TCP processing from the host system onto specialized processors which may have suitable software for handling TCP processing. The TCP offload engines may be configured to implement various TCP algorithms for handling faster network connections, thereby allowing host system processing resources to be allocated or reallocated to application processing.

[0006] In order to alleviate the consumption of host resources, a TCP connection can be offloaded from a host to a dedicated TCP/IP offload engine (TOE). Some of these host resources may include CPU cycles and subsystem memory bandwidth. During the offload process, TCP connection state information is offloaded from the host, for example from a host software stack, to the TOE. A TCP connection can be in any one of a plurality of states at a given time. To process the TCP connection, TCP software may be adapted to manage various TCP defined states. Being able to manage the various TCP defined states may require a high level of architectural complexity in the TOE.

[0007] Offloading state information utilized for processing a TCP connection to the TOE may not necessarily be the best solution because many of the states such as CLOSING, LAST\_ACK and FIN\_WAIT\_2 may not be performance sensitive. Furthermore, many of these non-performance sensitive states may consume substantial processing resources to handle, for example, error conditions and potentially malicious attacks. These are but some of the factors that substantially increase the cost of building and designing the TOE. In addition, a TOE that has control, transferred from the host, of all the state variables of a TCP connection may be quite complex, can use considerable processing power and may require and consume a lot of TOE on-board-memory. Moreover, the TCP connection offloaded to the TOE that has control, transferred from the host, of all the state variables of the TCP connection can be inflexible and susceptible to connection loss.

[0008] TCP segmentation is a technology that may permit a very small portion of TCP processing to be offloaded to a network interface card (NIC). In this regard, a NIC that supports TCP segmentation does not truly incorporate a full transmission control processing offload engine. Rather, a NIC that supports TCP segmentation only has the capability to segment outbound TCP blocks into packets having a size equivalent to that which the physical medium supports. Each of the outbound TCP blocks are smaller than a permissible TCP window size. For example, an Ethernet network interface card that supports TCP Segmentation, may segment a 4KB block of TCP data into 3 Ethernet packets. The maximum size of an Ethernet packet is 1518 bytes inclusive of header and a trailing CRC.

[0009] A device that supports TCP segmentation does track certain TCP state information such as the TCP sequence number that is related to the data that the offload NIC is segmenting. However, the device that

supports TCP segmentation does not track any state information that is related to inbound traffic, or any state information that is required to support TCP acknowledgements or flow control. A NIC that supports full TCP offload in the established state is responsible for handling TCP flow control, and responsible for handling incoming TCP acknowledgements, and generating outbound TCP acknowledgements for incoming data.

[0010] TCP segmentation may be viewed as a subset of TCP offload. TCP segmentation allows the protocol stack or operating system to pass information in the form of blocks of TCP data that has not been segmented into individual TCP packets to a device driver. The block of data may be greater than the size of an Ethernet packet. For instance, the block of data to be segmented could be 4 Kbytes or 16 Kbytes. A network adapter associated with the device driver may acquire the blocks of TCP data, packetize the acquired blocks of TCP data into 1518-byte Ethernet packets and update certain fields in each incrementally created packet. For example, the network adapter may update a corresponding TCP sequence number for each of the TCP packets by incrementing the TCP sequence number for each of the packets. In another example, an IP identification (IP ID) field and flag field would also have to be updated for each packet. One limitation with TCP segmentation is that TCP segmentation may only be done on a block of data that is less than a TCP window size. This is due to the fact that a device implementing TCP segmentation has no influence over TCP flow control. Accordingly, the device implementing TCP flow control only segment outbound TCP packets.

[0011] A TCP segmentation device does not examine incoming packets and as such, has no influence over flow control. Any received acknowledgement packet is passed up to the host for processing. In this regard, acknowledgement packets that are utilized for flow control are not processed by the TCP segmentation device. Moreover, a TCP segmentation device does not perform congestion control or "slow-start" and does not calculate or modify any variables that are passed back to the operating system and/or host system processor.

[0012] Another limitation with TCP segmentation is that information tracked by TCP segmentation is only information that is pertinent for the lifetime of the TCP data. In this regard, for example, the TCP segmentation device may track TCP segmentation numbers but not TCP acknowledgement (ACK) numbers. Accordingly, the TCP segmentation device tracks only a minimal subset of information related to corresponding TCP data. This limits the capability and/or functionality of the TCP segmentation device. A further limitation with TCP segmentation is that a TCP segmentation device does not pass TCP processed information back to an operating system and/or host processor. This lack of feedback limits the TCP processing that otherwise may be achieved by an operating system and/or host system processor.

[0013] Further limitations and disadvantages of con-

ventional and traditional approaches will become apparent to one of skill in the art, through comparison of such systems with some aspects of the present invention as set forth in the remainder of the present application with reference to the drawings.

## BRIEF SUMMARY OF THE INVENTION

[0014] Aspects of the invention may be found in, for example, systems and methods that provide TCP/IP offload. In one embodiment of the invention, a system for TCP/IP offload may include, for example, a host and a TCP/IP offload engine (TOE). The host may be coupled to the TOE. The host may transfer control of at least a portion of TCP connection variables associated with the TCP connection to the TOE. The TOE may update at least a portion of the TCP connection variables and transfer or feedback the updated TCP connection variables back to the host.

[0015] In accordance with another embodiment of the invention, a system is provided for TCP connection offload. The system may include, for example, a host and a network interface card (NIC) that may be coupled to the host. For a particular connection offloaded to the NIC, control of state information is split between the host and the NIC. Accordingly, information may be transferred to the NIC and the NIC may update at least a portion of the transferred information. Subsequently, the NIC may transfer at least a portion of the updated information back to the host where the host may utilize this information to manage this and/or another connection.

[0016] In another embodiment, the invention may provide a method for TCP/IP offload. The method may include, for example, one or more of the following: deciding to offload a particular TCP connection from a host to a TOE; transferring control of at least a portion of connection variables associated with the particular TCP connection from the host to the TOE; sending a snapshot of remaining connection variables whose control was not transferred to the TOE; and managing the particular TCP connection via the TOE using the connection variables transferred to the TOE and/or using the snapshot. At least a portion of updated connection variables and/or snapshot variables associated with the TCP connection may be transferred back to the host for processing by the host.

[0017] Another embodiment of TCP/IP offload method may include, for example, one or more of the following: deciding to offload an established TCP connection from a host to a TOE; transferring control of segment-variant variables to the TOE from the host; sending a snapshot of segment-invariant variables and connection-invariant variables to the TOE; and independently processing incoming TCP packets via the TOE based upon the segment-variant variables and the snapshot. The TOE may update at least a portion of the segment-variant variables and snapshot and transfer at least portions of the segment-variant variables and the snapshot

back to the host. In an embodiment of the invention, the host may handle all TCP states except possibly for the ESTABLISHED state which may be offloaded to the TOE.

[0018] The invention may also include a method that processes a TCP connection, which may include, for example, one or more of the following: establishing the TCP connection; sharing a control plane for the TCP connection between a host and a TOE; and communicating updated TCP connection variables from the TOE back to the host. Accordingly, at least a portion of the updated TCP connection variables may be utilized to control the TCP connection and/or another TCP connection.

[0019] In another embodiment of the invention, a method for TCP offload may include acquiring TCP connection variables from a host and managing at least one TCP connection using the acquired TCP connection variables. At least a portion of the acquired TCP connection variables may be updated and at least some of the updated TCP connection variables may be transferred back to the host. The TCP connection variables may be independent of bandwidth delay product. At least a portion of the updated TCP connection variables may be utilized by the host to process the TCP connection or another TCP connection. A stack may be utilized to transfer the TCP connection variables between at least the host and a TOE. In this regard, the TOE may pull the TCP connection variables from the stack and the host may push the TCP connection variables onto the stack. Also, the updated TCP connection variables may be placed on the stack by the TOE and the host may subsequently pull the updated TCP connection variables from the stack.

[0020] The invention may also provide a machine-readable storage, having stored thereon, a computer program having at least one code section for providing TCP offload. The at least one code section may be executable by a machine for causing the machine to perform steps which may include acquiring TCP connection variables from a host and managing at least one TCP connection using the acquired TCP connection variables. At least a portion of the acquired TCP connection variables may be updated and transferred back to the host. The TCP connection variables may be independent of bandwidth delay product. The machine-readable storage may further include code for utilizing at least a portion of the updated TCP connection variables to process the TCP connection or another TCP connection. In another aspect of the invention, the machine-readable storage may include code for pulling the TCP connection variables from a stack, code for pushing updated TCP connection variables onto the stack, and code for pulling connection variables from the stack.

According to an aspect of the invention, a system for providing TCP/IP offload is provided, comprising:

a host; and

a TCP/IP offload engine (TOE) coupled to said host, wherein said host transfers control of at least a portion of TCP connection variables to said TOE.

According to an aspect of the invention, a system for providing TCP/IP offload comprises:

a host; and  
a TCP/IP offload engine (TOE) coupled to said host, wherein said host transfers control of at least a portion of TCP connection variables to said TOE and said TOE provides updated TCP variables back to said host.

Advantageously, said host transfers control of segment-variant TCP connection variables to said TOE and said TOE provides updated TCP segment-variant TCP connection variables back to said host.

Advantageously, said TCP segment-variant TCP connection variables further comprises:

IP packet identifier;  
time remaining for retransmission;  
time remaining for delay acknowledgement;  
time remaining for keep alive;  
congestion window variables comprising:

congestion window (SND\_CWIN); and  
slow start threshold (SSTHRESH);

round trip time variables comprising:

smoothed round trip time (RTT);  
smoothed delta (DELTA); and  
time remaining for PUSH; and

TCP state and timestamp send and receive sequence variables comprising:

sequence number for first un-ACK'd data (SND\_UNA);  
sequence number for next send (SND\_NXT);  
maximum sequence number ever sent (SND\_MAX);  
maximum send window (MAX\_WIN);  
sequence number for next receive (RCV\_NXT); and  
receive window size (RCV\_WND).

Advantageously, said host provides said TOE with a snapshot of at least one of connection-invariant TCP connection variables and segment-invariant TCP connection variables, said TOE transferring updated at least one of connection-invariant TCP connection variables and segment-invariant TCP connection variables back to said host; and

wherein said TOE utilizes transferred segment-variant TCP connection variables and said snapshot of

said at least one of said connection-invariant TCP connection variables and said segment-invariant TCP connection variables to process at least one of incoming and outgoing TCP segments.

Advantageously, said TOE utilizes said transferred segment-variant TCP connection variables and a snapshot of at least one of connection-invariant TCP connection variables and segment-invariant TCP connection variables to independently process incoming TCP segments.

Advantageously, said host provides said TOE with a snapshot of connection-invariant TCP connection variables and segment-invariant TCP connection variables.

Advantageously, said host provides said TOE with a snapshot of TCP connection variables that are not segment-variant TCP connection variables.

Advantageously, said host posts at least one buffer that is utilized by any TCP connection, a buffer that is dedicated to one or more TCP connections and an application buffer.

Advantageously, said TOE manages said TCP connection.

Advantageously, said TOE manages at least one of segmentation, acknowledgement processing, windowing and congestion avoidance.

Advantageously, said TOE maintains exclusive read-write access to offloaded segment-variant variables.

Advantageously, said TOE exclusively updates offloaded segment-variant variables.

Advantageously, said host maintains read-write access to segment-invariant variables.

Advantageously, said TOE has read-only access to segment-invariant variables.

Advantageously, said host sends a message to said TOE concerning a change in a particular TCP connection variable whose control was not transferred to said TOE, and

said TOE updates said particular TCP connection variable.

Advantageously, said TCP connection is an ESTABLISHED state.

Advantageously, said TCP connection variables are TCP connection variables that are independent of bandwidth delay product.

Advantageously, said host provides connection setups.

Advantageously, the system provides resistance to DoS attacks by allowing said host to handle said connection setups.

Advantageously, said host handles all TCP states exclusive of an ESTABLISHED state which may be offloaded to said TOE.

Advantageously, said TOE handles only connections that are in performance sensitive states.

Advantageously, said host processes resource utilization statistics in helping to determine which con-

nections to offload and which connections to upload.

Advantageously, said host determines which connections to offload and which connections to upload.

Advantageously, at least one of said TOE and a device driver software for said TOE determines at least one of TCP connections to offload and TCP connections to upload.

According to an aspect of the invention, a system for providing connection offload comprises:

a host; and  
a network interface card (NIC) coupled to said host, wherein, for a particular connection offloaded to said NIC, control of state information is split between said host and said NIC and said NIC uploads at least a portion of updated connection variables for said particular connection to said host.

Advantageously, said particular connection employs a connection-oriented transport layer protocol (TLP).

Advantageously, said connection-oriented TLP comprises a TCP.

Advantageously, said host transfers control of segment variant variables corresponding to said particular connection to said NIC.

According to an aspect of the invention, a method for providing TCP/IP offload comprises:

deciding to offload a particular TCP connection from a host to a TOE;  
transferring control of connection variables of said particular TCP connection from said host to said TOE and transferring a snapshot of remaining connection variables whose control was not transferred to said TOE;  
managing said particular TCP connection via said TOE using said at least a portion of said connection variables transferred to said TOE and at least a portion of said snapshot; and  
updating at least a portion of said connection variables and a portion of said snapshot and transferring said updated at least said portion of said connection variables and said portion of said snapshot back to said host.

Advantageously, said one or more connection variables of said particular TCP connection transferred to said TOE comprise at least one segment-variant variables of said particular TCP connection.

Advantageously, said connection variables of said particular TCP connection transferred to said TOE lacks segment-invariant variables of said particular TCP connection.

Advantageously, said connection variables of said particular TCP connection transferred to the TOE lacks segment-invariant variables and connection-invariant variables of said particular TCP connection.

Advantageously, the method further comprises:

determining if at least one of said connection variables controlled by said host have changed;  
 notifying said TOE of changes in said at least one connection variables controlled by said host that has changed; and  
 updating said connection variables in said TOE in accordance with said notified changes.

According to an aspect of the invention, a method for providing TCP/IP offload comprises:

deciding to offload an established TCP connection from a host to a TOE;  
 transferring control of segment-variant variables to said TOE from said host;  
 sending a snapshot of segment-invariant variables and connection-invariant variables to said TOE;  
 independently processing incoming TCP packets via said TOE based upon said segment-variant variables and said snapshot; and  
 updating at least a portion of said sent snapshot and at least a portion of said segment-variant variables and transferring at least a portion of said updated at least said portion of said sent snapshot and at least said portion of said updated segment-variant variables back to said host.

According to an aspect of the invention, a method for processing a TCP connection comprises:

establishing the TCP connection; and  
 sharing a control plane for said TCP connection between a host and a TOE; and  
 communicating updated TCP connection variables from said TOE back to said host.

Advantageously, said sharing of said control plane comprises transferring control of segment-variant variables corresponding to said TCP connection to said TOE.

Advantageously, the method further comprises uploading said TCP connection to said host from said TOE.

Advantageously, uploading said TCP connection comprises transferring control of segment-variant variables corresponding to said TCP connection to said host.

Advantageously, the method further comprises offloading said uploaded TCP connection to said TOE from said host.

Advantageously, offloading said uploaded TCP connection comprises transferring said control of said segment-variant variables corresponding to said uploaded TCP connection to said TOE.

According to an aspect of the invention, a method for TCP offload comprises:

acquiring TCP connection variables from a host;  
 managing at least one TCP connection using said acquired TCP connection variables;  
 updating at least a portion of said acquired TCP connection variables; and  
 transferring said updated at least a portion of said acquired TCP connection variables back to said host.

Advantageously, said TCP connection variables are independent of bandwidth delay product.

Advantageously, the method further utilizes at least a portion of said updated at least said portion of said acquired TCP connection variables to process said at least said at least one TCP connection by said host.

Advantageously, the method further comprises pulling said TCP connection variables from a stack.

Advantageously, the method further comprises pushing said updated at least a portion of said acquired TCP connection variables onto a stack.

According to an aspect of the invention, a machine-readable storage is provided, having stored thereon, a computer program having at least one code section for providing TCP offload, the at least one code section being executable by a machine for causing the machine to perform steps comprising:

acquiring TCP connection variables from a host;  
 managing at least one TCP connection using said acquired TCP connection variables;  
 updating at least a portion of said acquired TCP connection variables; and  
 transferring said updated at least a portion of said acquired TCP connection variables back to said host.

Advantageously, said TCP connection variables are independent of bandwidth delay product.

Advantageously, the machine-readable storage further comprises code for utilizing at least a portion of said updated at least said portion of said acquired TCP connection variables to process said at least said at least one TCP connection by said host.

Advantageously, the machine-readable storage further comprises code for pulling said TCP connection variables from a stack.

Advantageously, the machine-readable storage further comprises code for pushing said updated at least a portion of said acquired TCP connection variables onto a stack.

[0021] These and other advantages, aspects and novel features of the present invention, as well as details of an illustrated embodiment thereof, will be more fully understood from the following description and drawings.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0022] FIG. 1 is a block diagram of a system that pro-

vides TCP/IP offload in accordance with an embodiment of the invention.

**[0023]** FIG. 2 is a flow chart illustrating exemplary steps for TCP/IP offloading in accordance with an embodiment of the invention.

**[0024]** FIG. 3 is a flow chart illustrating exemplary steps for providing TCP/IP offload in accordance with an embodiment of the invention.

**[0025]** FIG. 4 is a flow chart illustrating exemplary steps that may be utilized for TCP offload in accordance with an embodiment of the invention.

## DETAILED DESCRIPTION OF THE INVENTION

**[0026]** Certain aspects of the invention may provide a method for TCP offload, which may include acquiring TCP connection variables from a host and managing at least one TCP connection using the acquired TCP connection variables. At least a portion of the acquired TCP connection variables may be updated and at least some of the updated TCP connection variables may be transferred back to the host. In accordance with an aspect of the invention, the TCP connection variables may be variables that are independent of bandwidth delay product. At least a portion of the updated TCP connection variables may be utilized by the host to process the TCP connection or another TCP connection. A stack may be utilized to transfer the TCP connection variables between at least the host and a TOE. In this regard, the host may push the TCP connection variables onto the stack and the TOE may pull the TCP connection variables from the stack. Also, the updated TCP connection variables may be placed on the stack by the TOE and the host may subsequently pull the updated TCP connection variables from the stack.

**[0027]** With regard to TCP segmentation, each of the outbound TCP blocks are smaller than a permissible TCP window size utilized for TCP segmentation. However, the invention is not limited in this regard. Accordingly, in an aspect of the invention, a TOE device may have the capability to provide much further TCP processing and offload than a device that simply supports TCP segmentation. Various aspects of the invention may overcome the TCP segmentation limitation in which TCP segmentation may only be done on a block of data that is less than a TCP window size. In this regard, in order to overcome this limitation, in accordance with an aspect of the invention, since the TOE supports management of TCP flow control, the TOE may be adapted to segment large blocks of data down to the individual packets. The TOE may ensure that transmissions are scheduled such that the sender never sent data beyond the TCP window. Additionally, packetization in accordance with an embodiment of the invention may be done beyond the TCP window size. The TOE takes incoming received packets that are acknowledgment packets for the outbound TCP data stream and acknowledges those outbound packets. If the acknowl-

edgement packet causes the window size to increase, then more packets may be sent out by the TOE device in accordance with an aspect of the invention.

**[0028]** Although TCP segmentation is a transmit-only related technology that does limited TCP processing of transmitted packets, the TOE in accordance with various embodiments of the invention is not so limited. In this regard, the TOE in accordance with an embodiment of the invention may process and manage both transmitted and received packets. Furthermore, a much broader range of TCP processing and management may be done by the TOE in accordance with the invention than with a TCP segmentation device. For example, with TOE, TCP information may be passed to a NIC from an operating system and/or host system processor in such a manner that the NIC may be viewed as the owner of the TCP connection. The NIC may then manage and update the TCP state information, which may include but is not limited to, TCP segment numbers and acknowledgment numbers. Subsequent to the processing and/or updating of the TCP state information, the processed and/or updated information may be passed back to an operating system and/or host system processor. The host or system processor may then utilize the information passed back to it from the NIC. Notably, TCP segmentation does not provide this feedback of information to the host system processor and/or operating system.

**[0029]** Certain embodiments of the invention may also provide a robust and efficient transmission control protocol/internet protocol (TCP/IP) offload scheme that may be adapted, for example, to allow the partition of TCP processing between a TCP/IP offload engine (TOE) and a host TCP/IP implementation. The host TCP/IP implementation may include one or more host TCP/IP applications and one or more host processors. For example, in one aspect of the invention, the TCP offload scheme may offload the connections that are in an ESTABLISHED state to the TOE. In other words, aspects of the invention may include the offloading of corresponding TCP state variables that may be utilized, for example, during the ESTABLISHED state. Accordingly, the TCP/IP offload scheme may split a TCP control plane between the host software and the TOE. The TOE may be designed, for example, to implement a subset or a minimum subset of the TCP control plane which may be less complex to implement and may utilize less memory. The TOE, which may be adapted to such an offload scheme, may be implemented in a cost effective manner. The more complicated aspects of TCP connection management may be handled, for example, by the host software and may provide greater reliability and flexibility.

**[0030]** FIG. 1 is a block diagram of a system that provides TCP/IP offload in accordance with an embodiment of the invention. Referring to FIG. 1, the system may include, for example, a host 10, host application software 12 and a TOE 20. The host 10 may include, for

example, a host CPU 30 and a host memory 40. The host memory 40 may be adapted to include, for example, an application buffer 50. The application buffer 50 may be adapted to include, for example, a transmission application buffer (TxBuf) 60 and a receive application buffer (RxBuf) 70. The TOE 20 may include, for example, a direct memory access (DMA) engine 25 and a FIFO buffer 70.

**[0031]** The host 10 may be coupled to the TOE 20 via a host interface 80. The host interface may include, but is not limited to a peripheral component interconnect (PCI) bus, PCI-X bus, ISA, SCSI or any other suitable bus. The TOE 20 may be coupled to a physical communications medium 90. The physical communication medium 90 may be a wired medium, wireless medium or a combination thereof. The physical communication medium 90 may include, but is not limited to, Ethernet and fibre channel. Although illustrated on opposite sides of the host interface 80, the host 10 may be, at least in part, disposed on a network interface card (NIC) that includes the TOE 20. Accordingly, in an aspect of the invention, the TCP state plane may be split between the host 10 and the TOE 20.

**[0032]** In one embodiment, a TCP connection may be completely described, for example, by three different sets of variables. The three sets of variables may be, for example, connection-invariant variables, segment-invariant variables and segment-variant variables. The connection-invariant variables may be constant during the lifetime of the TCP connection. The segment-invariant variables may not change from TCP segment to TCP segment, but may change from time to time during the lifetime of the TCP connection. The segment-variant variables may change from TCP segment to TCP segment.

**[0033]** Connection-invariant variables may include, for example, source IP address, destination IP address, IP time-to-live (TTL), IP type-of-service (TOS), source TCP port number, destination TCP port number, initial send sequence number, initial receive sequence number, send window scaling factor and receive window scaling factor.

**[0034]** Segment-invariant variables may include, but are not limited to, source MAC address, next hop's MAC address, MAC layer encapsulation, effective maximum segment size, keep-alive intervals and maximum allowance and flags such as, for example, nagle algorithm enable and keep-alive enable.

**[0035]** Segment-variant variables may include, but are not limited to, IP packet identifier; send and receive sequence variables such as, for example, sequence number for first un-acked data (SND\_UNA), sequence number for next send (SND\_NXT), maximum sequence number ever sent (SND\_MAX), maximum send window (MAX\_WIN), sequence number for next receive (RCV\_NXT) and receive window size (RCV\_WND). Additional exemplary segment-variant variables may include congestion window variables such as congestion

window (SND\_CWIN) and slow start threshold (SS-THRESH) round trip time variables which may include, but are not limited to, smoothed round trip time (RTT) and smoothed delta (DELTA). Other exemplary segment-variant variables may include time remaining for retransmission, time remaining for delay acknowledgment, time remaining for keep alive, time remaining for PUSH and TCP state and timestamp.

**[0036]** During operation, if a TCP connection is not offloaded, then at least some of the three sets of variables including the connection-invariant variables, the segment-invariant variables and the segment-variant variables may be owned by the host software of the host 10. If the TCP connection is not offloaded, then the TOE 20 may not have access to these variables. However, once the variables are offloaded, the TOE 20 may be configured to update the variables which may be associated with both transmission and reception and pass the updated transmission and reception variables back to the host 10. In this regard, the TOE may update variables that are independent of TCP delay bandwidth product and pass these updated variables back to the host 10 for processing.

**[0037]** FIG. 2 is a flow chart illustrating exemplary steps for TCP/IP offloading in accordance with an embodiment of the invention. Referring to FIG. 2, if a connection is offloaded to the TOE 20, then in step 202, the host software may transfer control of the segment-variant variables to the TOE 20. In one example, a portion of the host software protocol control block or TCP control block may be transferred to the TOE 20. In step 204, the host software may take a snapshot of the remaining variables such as the connection-invariant variables and/or the segment invariant variables and send the snapshot to the TOE 20. In one example, the snapshot may be used over and over again by the TOE 20. In step 206, the host software may post a buffer in the host memory 40. For example, the host software may post the application buffer 50 in the host memory 40 and may set up the transmit application buffer (TxBuf) 60 and the receive application buffer (RxBuf) 70 in the application buffer 50. In step 208, the TOE 20 may be responsible for managing the complete TCP connection, including, for example, segmentation, acknowledgement processing, windowing and congestion avoidance. In step 210, at least a portion of the variables that have been updated may be transferred back to the host for processing.

**[0038]** For example, by controlling the segment-variant variables and using the snapshot of the remaining variables, the TOE 20 may process or independently process, incoming TCP segments from the physical communications medium 90 and may place at least a portion such as a payload, of the incoming TCP segments into the host memory 40 via the DMA engine 25. In this regard the incoming TCP segment payload may be placed in the RX application buffer 70 portion of the application buffer 50 via the DMA engine 25.

**[0039]** In one embodiment of the invention, while the



TOE 20 may be adapted to manage the complete TCP connection, the TOE 20 may have exclusive read-write access to offloaded segment-variant variables and may exclusively update the offloaded segment-variant variables. The host software or host application software 12 may have read-write access to the segment-invariant variables. The TOE 20 may have read-only access to the segment-invariant variables. If the host application software 12 changes the variables such as the next hop's MAC address, the host application software 12 may notify the TOE 20 by, for example, sending a message to the TOE 20. The TOE 20 may then update the variables. The updated variables may be fed back to the host application software 12 where they may be utilized for TCP processing, for example. Accordingly, the connection-invariant variables may exist in both the host software and the TOE 20.

[0040] FIG. 3 is a flow chart illustrating exemplary steps for providing TCP/IP offload in accordance with an embodiment of the invention. Referring to FIG. 3, in step 302, the host 10 may determine whether one or more of the connection variables such as the segment-invariant variables controlled by the host 10 have changed. For example, the host software may change one or more variables such as a next hop MAC address. If one or more of the connection variables controlled by the host 10 are not changed, then the process may be complete. If one or more of the connection variables controlled by the host 10 are changed, then, in step 304, the host software may notify the TOE 20 of the change in the one or more connection variables controlled by the host 10. In step 306, the TOE 20 may accordingly update one or more of the variables. In step 308, the TOE may pass the updated variables back to the host 10.

[0041] Some embodiments according to the present invention may include one or more of the following advantages. Some embodiments may be more reliable and may provide for the uploading of connection from the TOE to the host and offloading of connections from the host to the TOE at any time. Since less state information may be kept by the TOE hardware, uploading and offloading, for example, selected connections can be accelerated. An offloaded connection may be uploaded by returning control of, for example, the segment-variant variables corresponding to the offloaded connection back to the host 10. The uploaded connection may subsequently be offloaded by transferring, for example, the control of the segment-variant variables corresponding to the uploaded connection to the TOE 20.

[0042] FIG. 4 is a flow chart illustrating exemplary steps that may be utilized for TCP offload in accordance with an embodiment of the invention. Referring to FIG. 4, in step 402, a TOE may acquire or receive variables that are independent of the bandwidth delay product from a host system. In step 404, the TOE may manage the connection utilizing the acquired or received variables

that are independent of the bandwidth delay product. In step 406, the TOE may update at least a portion of the acquired variables that are independent of the bandwidth delay product. In step 408, at least a portion of the updated variables that are independent of the bandwidth may be transferred back to the host. In step 410, the host may utilize the updated variables that are independent of the bandwidth delay product that have been transferred to it for TCP processing.

[0043] In accordance with an aspect of the invention, a stack 14 may be utilized to facilitate the transfer of the variables that are independent of the bandwidth delay product. The stack 14 may be implemented in hardware, software or a combination thereof. Notwithstanding, the TOE may be adapted to pull information from the stack 14 and to push updated information onto the stack 14. The host may also be adapted to push TCP information onto the stack 14 and to pull the updated information from the stack 14. Accordingly, with reference to step 402, the TOE may pull the variables that are independent of the bandwidth delay product from the stack 14. With reference to step 406, after the TOE updates the acquired variables that are independent of the bandwidth delay product, the updated variables that are independent of the bandwidth delay product may be pushed onto the stack 14. In this regard, with reference to step 408, the host may then pull the updated variables that are independent of the bandwidth delay product from the stack 14.

[0044] The TOE may provide a more flexible approach to TCP processing compared to a TCP Segmentation offload device, since the TOE device may facilitate TCP processing on both the received side and the transmit side. Additionally, since the TOE may be adapted to handle receive and transmit variables, the TOE provides a more flexible and efficient methodology for supporting the efficient setup and tear down of network connections.

[0045] Certain embodiments of the invention may offer better resistance against denial-of-service (DoS) attacks or other attacks as connection setup may be handled by a host that is more flexible and more powerful than the TOE NIC. In a DoS attack, an attacker attempts to consume as many resources on the targeted or attacked system, thereby preventing the targeted system from providing services to other network devices. The frequent introduction of new attacks may make a flexible host with sufficient memory and CPU power a better choice for running connection setup. The flexible host may be a better choice than, for example, a particular hardware TOE that may have limited code space, computer power, system knowledge and flexibility. In addition, the decision to honor a connection request may, at times, be based upon, for example, sophisticated and dynamic heuristics.

[0046] Aspects of the invention may also provide better overall system performance and efficiency. The TOE NIC may be more efficient in handling, for example, con-

nections that are in performance sensitive states of the TCP state machine. In particular, when the TOE NIC handles only connections that are in performance sensitive states of the TCP state machine, additional limited hardware resources may become available. Accordingly, the TOE NIC may be adapted to upload connections that are no longer in performance sensitive states and to offload connections that are in performance sensitive states. Such actions may positively impact such figures of merit such as, for example, hardware TOE efficiency. Other aspects of the invention may be more efficient and may provide better over all system performance because, for example, the host may use flexible, changing, easy-to-update, easy-to-upgrade and more sophisticated algorithms to decide which connections to offload or to upload.

[0047] Some embodiments according to the present invention may provide statistics to the host relating to resource utilization. The statistics may include, for example, one or more of the following: available resources; utilization of bandwidth per offloaded connection; number of frames per offloaded connection; errors per offloaded connection; change of state of a transport layer protocol (TLP) such as, for example, TCP, or an upper layer protocol (ULP); trend of utilization such as uptake in rate, slow down, for example; and resource consumption per offloaded connection. The host may use the statistical information at its own discretion to help drive the upload or offload decision process. For example, the host may utilize the statistical information to upload some connections while offloading others. The host may also contemplate other criteria such as modes of operation, computation or network load profiles, presently executed applications and roles in the network, for example. Some of these criteria may be dynamic criteria.

[0048] Certain embodiments of the invention may also provide fail-over support from a failed TOE NIC to an operating TOE NIC. Fail-over may include, for example, designating a NIC as having failed when the network cable is unplugged from the network or any other failure of an existing network link. Thus, even though the hardware of one TOE NIC may fail, the connection may still be maintained by transferring state information associated with the failed TOE NIC to another functional TOE NIC. The robustness of the transfer may be further enhanced by part of the connection state information being maintained by the host and part of the connection state information being maintained by the TOE NIC.

[0049] Accordingly, the present invention may be realized in hardware, software, or a combination of hardware and software. The present invention may be realized in a centralized fashion in one computer system or in a distributed fashion where different elements are spread across several interconnected computer systems. Any kind of computer system or other apparatus adapted for carrying out the methods described herein is suited. A typical combination of hardware and software may be a general-purpose computer system with

a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein.

[0050] The present invention also may be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which when loaded in a computer system is able to carry out these methods. Computer program in the present context means any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following: a) conversion to another language, code or notation; b) reproduction in a different material form.

[0051] While the present invention has been described with reference to certain embodiments, it will be understood by those skilled in the art that various changes may be made and equivalents may be substituted without departing from the scope of the present invention. In addition, many modifications may be made to adapt a particular situation or material to the teachings of the present invention without departing from its scope. Therefore, it is intended that the present invention not be limited to the particular embodiment disclosed, but that the present invention will include all embodiments falling within the scope of the appended claims.

## Claims

### 1. A system for providing TCP/IP offload, comprising:

a host; and  
a TCP/IP offload engine (TOE) coupled to said host, wherein said host transfers control of at least a portion of TCP connection variables to said TOE.

### 2. The system according to claim 1, wherein said TOE provides updated TCP variables back to said host

### 3. The system according to claim 1 or 2, wherein said host transfers control of segment-variant TCP connection variables to said TOE and said TOE provides updated TCP segment-variant TCP connection variables back to said host.

### 4. A system for providing connection offload, comprising:

a host; and  
a network interface card (NIC) coupled to said host, wherein, for a particular connection offloaded to said NIC, control of state information is split between said host and said NIC and said

NIC uploads at least a portion of updated connection variables for said particular connection to said host.

5. A method for providing TCP/IP offload, comprising: 5

deciding to offload a particular TCP connection from a host to a TOE;  
transferring control of connection variables of said particular TCP connection from said host to said TOE and transferring a snapshot of remaining connection variables whose control was not transferred to said TOE;  
managing said particular TCP connection via said TOE using said at least a portion of said connection variables transferred to said TOE and at least a portion of said snapshot; and  
updating at least a portion of said connection variables and a portion of said snapshot and transferring said updated at least said portion of said connection variables and said portion of said snapshot back to said host. 10 15 20

6. The method according to claim 5, wherein said one or more connection variables of said particular TCP connection transferred to said TOE comprise at least one segment-variant variables of said particular TCP connection. 25

7. A method for providing TCP/IP offload, comprising: 30

deciding to offload an established TCP connection from a host to a TOE; transferring control of segment-variant variables to said TOE from said host; sending a snapshot of segment-invariant variables and connection-invariant variables to said TOE;  
independently processing incoming TCP packets via said TOE based upon said segment-variant variables and said snapshot; and  
updating at least a portion of said sent snapshot and at least a portion of said segment-variant variables and transferring at least a portion of said updated at least said portion of said sent snapshot and at least said portion of said updated segment-variant variables back to said host. 35 40 45

8. A method for processing a TCP connection, comprising: 50

establishing the TCP connection; and  
sharing a control plane for said TCP connection between a host and a TOE; and  
communicating updated TCP connection variables from said TOE back to said host. 55

9. A method for TCP offload, the method comprising:

acquiring TCP connection variables from a host;  
managing at least one TCP connection using said acquired TCP connection variables;  
updating at least a portion of said acquired TCP connection variables; and  
transferring said updated at least a portion of said acquired TCP connection variables back to said host.

10. A machine-readable storage, having stored thereon, a computer program having at least one code section for providing TCP offload, the at least one code section being executable by a machine for causing the machine to perform steps comprising:

acquiring TCP connection variables from a host;  
managing at least one TCP connection using said acquired TCP connection variables;  
updating at least a portion of said acquired TCP connection variables; and  
transferring said updated at least a portion of said acquired TCP connection variables back to said host.




FIG. 1

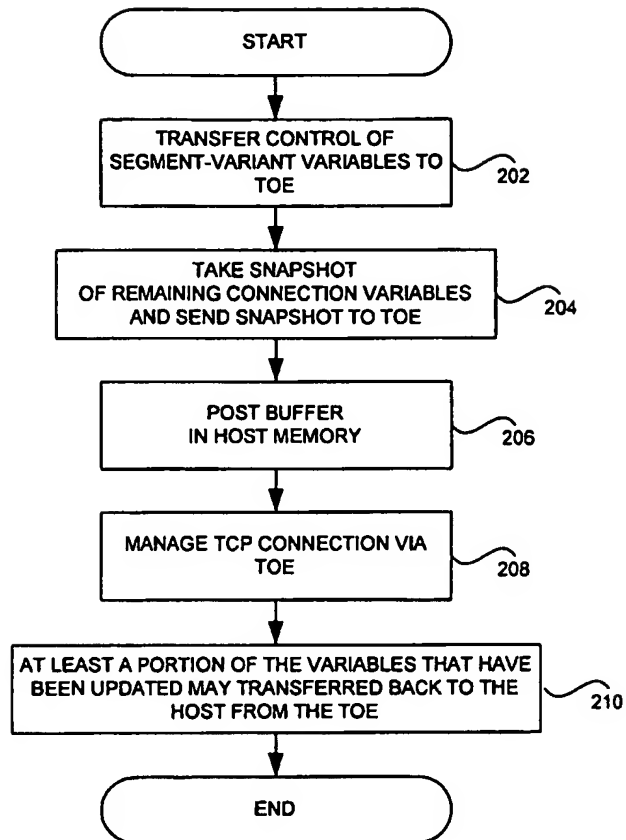


FIG. 2




FIG. 3

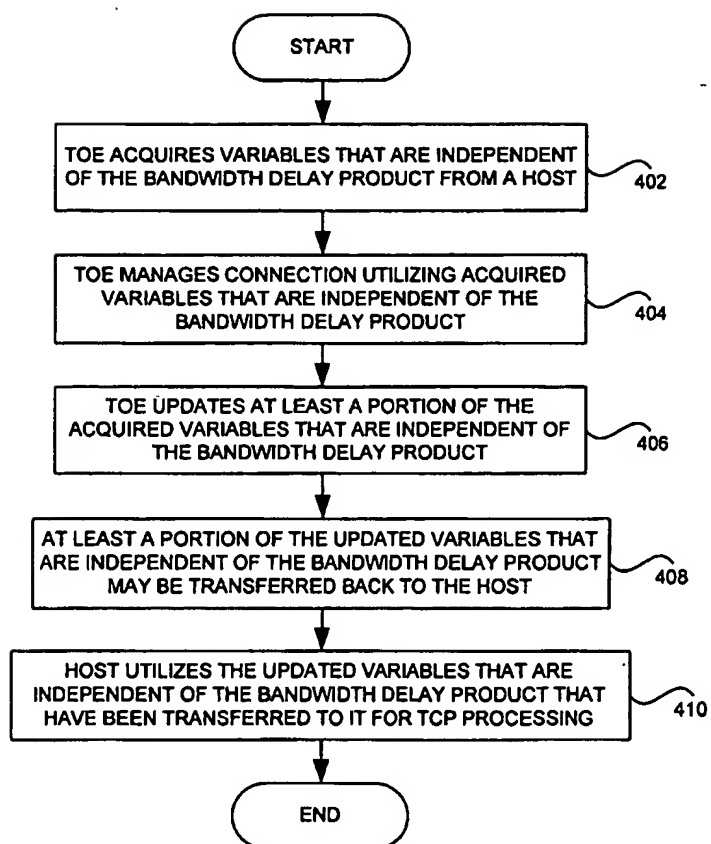


FIG. 4